

CLAIMS

1 1. A shared memory packet switching device having a plurality of receive ports for
2 receiving data packets, and a plurality of transmit ports for transmitting data packets, comprising:

3 a shared memory providing a shared memory space for temporary storage of data packets
4 received via the receive ports;

5 at least one input logic unit associated with at least one of the receive ports, and being
6 operative to determine whether said associated receive port is saturated by determining whether a
7 number of packets received via said associated receive port and currently stored in said shared
8 memory exceeds a predetermined drop threshold value;

9 a packet routing control unit communicatively coupled with said at least one input logic
10 unit, and being operative to determine a destination one of the transmit ports for each of the
11 received data packets; and

12 at least one output logic unit associated with at least one of the transmit ports, said output
13 logic unit being communicatively coupled with said packet routing control unit, and being
14 operative to determine whether said associated transmit port is congested by determining
15 whether a number of packets currently stored in said shared memory that are to be transmitted
16 via said associated transit port exceeds a predetermined congestion threshold value, and also
17 being operative to generate an associated output full signal indicative of whether said associated
18 transmit port is congested;

19 said input logic unit being responsive at least in part to each of said output full signals,
20 and being further operative to cause a selected packet received via said associated receive port to
21 be dropped if said associated receive port is currently saturated and said output full signals
22 indicate that a destination transmit port associated with said selected packet is currently
23 congested.

1 2. A shared memory packet switching device as recited in claim 1 further comprising:

2 a memory control unit for accessing data packets stored in said shared memory, and

3 being operative to generate a count enable signal, and to assert said enable signal while a
4 predetermined threshold portion of said shared memory space is occupied by stored packet data;

5 said input logic unit further including,

6 an input counter responsive to said enable signal, and operative to provide an
7 associated input count value indicative of the number of packets that are currently stored
8 in said shared memory and that have been received via said associated receive port during
9 a period wherein said enable signal is asserted, said input counter being further operative
10 to clear said associated input count value if said enable signal is de-asserted,

11 a first comparator unit responsive to said input count value, and operative to
12 generate an associated select drop signal, and to assert said associated select drop signal
13 based on a comparison between said input count value and said drop threshold value, and

14 a drop logic unit responsive to said associated select drop signal and at least in
15 part to said output full signals, and operative to drop a selected data packet received via
16 said associated receive port if said associated select drop signal is asserted and said
17 output full signals indicate that said destination transmit port associated with said selected
18 packet is currently congested.

1 3. A shared memory packet switching device as recited in claim 1 wherein:

2 said input logic unit further includes a second comparator unit responsive to said input
3 count value, and operative to generate an associated pause signal for indicating that backpressure
4 is to be asserted at said associated receive port, and to assert said associated pause signal based
5 on a comparison between said associated input count value and a predetermined backpressure
6 threshold value; and

7 said output logic unit further includes a transmit port control unit responsive to said
8 associated pause signal, and operative to assert back pressure on an associated network link that
9 is communicatively coupled with said associated receive port.

1 4. A shared memory packet switching device as recited in claim 2 wherein:

2 said input logic unit further includes a second comparator unit responsive to said input
3 count value, and operative to generate an associated pause signal for indicating that backpressure
4 is to be asserted at said associated receive port, and to assert said associated pause signal based
5 on a comparison between said associated input count value and a predetermined backpressure
6 threshold value; and

7 said output logic unit further includes a transmit port control unit responsive to said
8 associated pause signal, and operative to assert back pressure on an associated network link that
9 is communicatively coupled with said associated receive port.

1 5. A shared memory packet switching device as recited in claim 3 wherein said transmit
2 port control unit is operative to assert backpressure by transmitting a pause message via said
3 associated network link.

1 6. A shared memory packet switching device as recited in claim 1 wherein said output logic
2 unit further comprises:

3 an output counter operative to generate an associated output count value indicative of
4 said number of packets currently stored in said shared memory that are to be transmitted via said
5 associated transmit port; and

6 a comparator unit responsive to said output count value, and operative to generate said
7 associated output full signal based on a comparison between said output count value and said
8 predetermined congestion threshold value.

1 7. A shared memory packet switching device as recited in claim 2 wherein:

2 said packet routing unit is further operative to generate a plurality of transmit signals each
3 being associated with one of said transmit ports, and to assert a particular one of said transmit
4 signals when an associated packet is to be transmitted via said associated transmit port; and

5 said input logic unit is further responsive to said transmit signals, and
6 further operative to cause said selected packet to be dropped if said particular receive port is
7 currently saturated, said output full signals indicate that said destination transmit port associated
8 with said selected packet is currently congested, and said transmit signals indicate that said
9 selected packet is to be transmitted via said congested transmit port.

1 8. A shared memory packet switching device as recited in claim 7 wherein said input logic
2 unit is further responsive to a drop enable signal selectively enabled by a user of the switching
3 device to indicate whether packets may be dropped by the switching device, said input logic unit

4 only being operative to cause said selected packet to be dropped if said drop enable signal
5 indicates that packet dropping is enabled.

1 9. A shared memory packet switching device as recited in claim 2 wherein each of said
2 receive ports and an associated one of said transmit ports is formed by an associated one of N
3 bidirectional ports each being communicatively coupled with an associated network link.

1 10. A shared memory packet switching device as recited in claim 9 wherein said congestion
2 threshold value is proportional to a value determined by dividing said shared memory space into
3 N parts.

1 11. A shared memory packet switching device as recited in claim 9 wherein a maximum
2 portion of said shared memory space is allocated for packets received via each associated one of
3 the receive ports, and wherein said maximum portion is defined by dividing a reserved portion of
4 said shared memory space into N parts, said reserved portion being defined as the difference
5 between said shared memory space and said predetermined threshold portion.

1 12. A shared memory packet switching device as recited in claim 3 wherein said drop
2 threshold value is predefined to be greater than said backpressure threshold value, whereby said
3 switching device is operative to assert backpressure at a particular receive port before dropping
4 packets received via said particular receive port.

1 13. A shared memory packet switching device as recited in claim 3 wherein said drop
2 threshold value is defined to be less than said backpressure threshold value, whereby said
3 switching device is operative to drop packets received via said associated receive port in order to
4 avert the necessity of asserting backpressure at said associated receive port.

1 14. A shared memory packet switching device as recited in claim 11 wherein said drop
2 threshold value is substantially equal to a number of data packets which may be stored in said
3 maximum portion of said shared memory space.

1 15. A shared memory packet switching device as recited in claim 11 wherein the difference
2 between said maximum portion and said backpressure threshold value is greater than or equal to
3 N, whereby said switching device is operative to store at least one packet received at each of the
4 receive ports after asserting backpressure at said associated receive port.

1 16. A shared memory packet switching device having a plurality of receive ports for
2 receiving data packets, and a plurality of transmit ports for transmitting data packets, comprising:

3 a shared memory providing a shared memory space for temporary storage of data packets
4 received via the receive ports;

5 at least one input logic unit associated with at least one of the receive ports, and being
6 operative to determine whether said associated receive port is saturated by determining whether a
7 number of packets received via said associated receive port and currently stored in said shared
8 memory exceeds a predetermined drop threshold value;

9 a packet routing control unit communicatively coupled with said at least one input logic
10 unit, and being operative to determine a destination one of the transmit ports for each of the
11 received data packets, said packet routing unit being further operative to generate a plurality of
12 transmit signals each being associated with one of said transmit ports, and to assert a particular
13 one of said transmit signals when a received packet is to be transmitted via said associated
14 transmit port; and

15 at least one output logic unit associated with at least one of the transmit ports, said output
16 logic unit being communicatively coupled with said packet routing control unit, and being
17 operative to determine whether said associated transmit port is congested by determining
18 whether a number of packets currently stored in said shared memory that are to be transmitted
19 via said associated transit port exceeds a predetermined congestion threshold value, and also
20 being operative to generate an associated output full signal indicative of whether said associated
21 transmit port is congested;

22 said packet routing control unit also being responsive to said output full signals, and
23 being operative to generate a plurality of filter signals for indicating that a received packet is
24 destined for a congested one of the transmit ports;

25 said input logic unit being further responsive to each of said filter signals, and being
26 further operative to cause a selected packet received via said associated receive port to be

27 dropped if said associated receive port is currently saturated and said filter signals indicate that a
28 destination transmit port associated with said selected packet is currently congested.

1 17. A shared memory packet switching device as recited in claim 16 further comprising:
2 a memory control unit for accessing data packets stored in said shared memory, and
3 being operative to generate a count enable signal, and to assert said enable signal while a
4 predetermined threshold portion of said shared memory space is occupied by stored packet data;
5 said input logic unit further including,

6 an input counter responsive to said enable signal, and operative to provide an
7 associated input count value indicative of the number of packets that are currently stored
8 in said shared memory and that have been received via said associated receive port during
9 a period wherein said enable signal is asserted, said input counter being further operative
10 to clear said associated input count value if said enable signal is de-asserted,

11 a first comparator unit responsive to said input count value, and operative to
12 generate an associated select drop signal, and to assert said associated select drop signal
13 based on a comparison between said input count value and said drop threshold value, and

14 a drop logic unit responsive to said associated select drop signal and at least in
15 part to said output full signals, and operative to drop a selected data packet received via
16 said associated receive port if said associated select drop signal is asserted and said
17 output full signals indicate that said destination transmit port associated with said selected
18 packet is currently congested.

1 18. A shared memory packet switching device as recited in claim 16 wherein:
2 said input logic unit further includes a second comparator unit responsive to said input
3 count value, and operative to generate an associated pause signal for indicating that backpressure
4 is to be asserted at said associated receive port, and to assert said associated pause signal based
5 on a comparison between said associated input count value and a predetermined backpressure
6 threshold value; and
7 said output logic unit further includes a transmit port control unit responsive to said
8 associated pause signal, and operative to assert back pressure on an associated network link that
9 is communicatively coupled with said associated receive port.

1 19. A shared memory packet switching device as recited in claim 17 wherein:
2 said input logic unit further includes a second comparator unit responsive to said input
3 count value, and operative to generate an associated pause signal for indicating that backpressure
4 is to be asserted at said associated receive port, and to assert said associated pause signal based
5 on a comparison between said associated input count value and a predetermined backpressure
6 threshold value; and

7 said output logic unit further includes a transmit port control unit responsive to said
8 associated pause signal, and operative to assert back pressure on an associated network link that
9 is communicatively coupled with said associated receive port.

1 20. A shared memory packet switching device as recited in claim 18 wherein said transmit
2 port control unit is operative to assert backpressure by transmitting a pause message via said
3 associated network link.

1 21. A shared memory packet switching device as recited in claim 16 wherein said output
2 logic unit further comprises:

3 an output counter operative to generate an associated output count value indicative of
4 said number of packets currently stored in said shared memory that are to be transmitted via said
5 associated transmit port; and

6 a comparator unit responsive to said output count value, and operative to generate said
7 associated output full signal based on a comparison between said output count value and said
8 predetermined congestion threshold value.

1 22. A shared memory packet switching device as recited in claim 16 wherein said input logic
2 unit is further responsive to a drop enable signal selectively enabled by a user of the switching
3 device to indicate whether packets may be dropped by the switching device, said input logic unit
4 only being operative to cause said selected packet to be dropped if said drop enable signal
5 indicates that packet dropping is enabled.

1 23. A shared memory packet switching device as recited in claim 16 wherein each of said
2 receive ports and an associated one of said transmit ports is formed by an associated one of N
3 bidirectional ports each being communicatively coupled with an associated network link.

1 24. A shared memory packet switching device as recited in claim 23 wherein said congestion
2 threshold value is proportional to a value determined by dividing said shared memory space into
3 N parts.

1 25. A shared memory packet switching device as recited in claim 18 wherein said drop
2 threshold value is predefined to be greater than said backpressure threshold value, whereby said
3 switching device is operative to assert backpressure at a particular receive port before dropping
4 packets received via said particular receive port.

1 26. A shared memory packet switching device as recited in claim 18 wherein said drop
2 threshold value is defined to be less than said backpressure threshold value, whereby said
3 switching device is operative to drop packets received via said associated receive port in order to
4 avert the necessity of asserting backpressure at said associated receive port.

1 27. A process of controlling the flow of data through a shared memory packet switching
2 device having a plurality of receive ports for receiving data packets, a plurality of transmit ports
3 for transmitting data packets, and a shared memory providing a shared memory space for
4 temporary storage of data packets received via the receive ports, comprising the steps of:

5 (a) receiving a packet via an associated one of the receive ports;

6 (b) determining whether said associated receive port is currently saturated by determining
7 whether a number of packets received via said associated receive port and currently stored in the
8 shared memory exceeds a predetermined drop threshold value;

9 (c) determining a destination one of the transmit ports associated with said received data
10 packet;

11 (d) determining whether said destination transmit port is currently congested by
12 determining whether a number of packets currently stored in the shared memory that are to be

13 transmitted via said destination transmit port exceeds a predetermined congestion threshold
14 value; and

15 (e) dropping said received packet if said associated receive port is currently saturated and
16 said destination transmit port is currently congested.

1 28. A process as recited in claim 27 wherein said step (b) of determining whether said
2 associated receive port is currently saturated further comprises the steps of:

3 determining whether a currently occupied portion of the shared memory space is greater
4 than or equal to a predetermined threshold portion of the shared memory space;

5 if said occupied portion is greater than or equal to said threshold portion, enabling an
6 input counter to begin counting the number of packets received via said associated receive port
7 and currently stored in the shared memory; and

8 if said occupied portion is not greater than or equal to said threshold portion, resetting
9 said input counter.

1 29. A process as recited in claim 28 wherein said step (b) of determining whether said
2 associated receive port is currently saturated further comprises the steps of:

3 if said occupied portion is greater than or equal to said threshold portion,

4 increasing an input count value upon receiving a packet via said associated
5 receive port, and

6 decreasing said input count value upon transmitting a packet that has been
7 received via said associated receive port, via one of the transmit ports.

1 30. A process as recited in claim 27 further comprising the steps of:

2 (f) determining whether said number of packets received via said associated receive port
3 and currently stored in the shared memory exceeds a predetermined backpressure threshold
4 value; and

5 (g) if said number of packets received via said associated receive port and currently
6 stored in the shared memory exceeds said backpressure threshold value, asserting backpressure
7 at said associated receive port.

1 31. A process as recited in claim 28 further comprising the steps of:
2 (f) determining whether said number of packets received via said associated receive port
3 and currently stored in the shared memory exceeds a predetermined backpressure threshold
4 value; and
5 (g) if said number of packets received via said associated receive port and currently
6 stored in the shared memory exceeds said backpressure threshold value, asserting backpressure
7 at said associated receive port.

1 32. A process as recited in claim 30 wherein said step (g) further includes asserting
2 backpressure by transmitting a pause message via a network link that is communicatively
3 coupled with said associated receive port.

1 33. A process as recited in claim 30 wherein each of said receive ports and an associated one
2 of said transmit ports is formed by an associated one of N bidirectional ports each being
3 communicatively coupled with an associated network link.

1 34. A process as recited in claim 31 wherein each of said receive ports and an associated one
2 of said transmit ports is formed by an associated one of N bidirectional ports each being
3 communicatively coupled with an associated network link.

1 35. A process as recited in claim 33 wherein said congestion threshold value is proportional
2 to a value determined by dividing said shared memory space into N parts.

1 36. A process as recited in claim 34 wherein a maximum portion of said shared memory
2 space is allocated for packets received via each associated one of the receive ports, and wherein
3 said maximum portion is defined by dividing a reserved portion of said shared memory space
4 into N parts, said reserved portion being defined as the difference between said shared memory
5 space and said predetermined threshold portion.

1 37. A process as recited in claim 30 wherein said drop threshold value is predefined to be
2 greater than said backpressure threshold value, whereby said switching device is operative to

3 assert backpressure at a particular receive port before dropping packets received via said
4 particular receive port.

1 38. A process as recited in claim 30 wherein said drop threshold value is defined to be less
2 than said backpressure threshold value, whereby said switching device is operative to drop
3 packets received via said associated receive port in order to avert the necessity of asserting
4 backpressure at said associated receive port.

1 39. A process as recited in claim 36 wherein said drop threshold value is substantially equal
2 to a number of data packets which may be stored in said maximum portion of said shared
3 memory space.

1 40. A process as recited in claim 36 wherein the difference between said maximum portion
2 and said backpressure threshold value is greater than or equal to N, whereby said switching
3 device is operative to store at least one packet received at each of the receive ports after asserting
4 backpressure at said associated receive port.
5